WHITE PAPER

Isolating Microsoft Active Directory with Micro-Segmentation

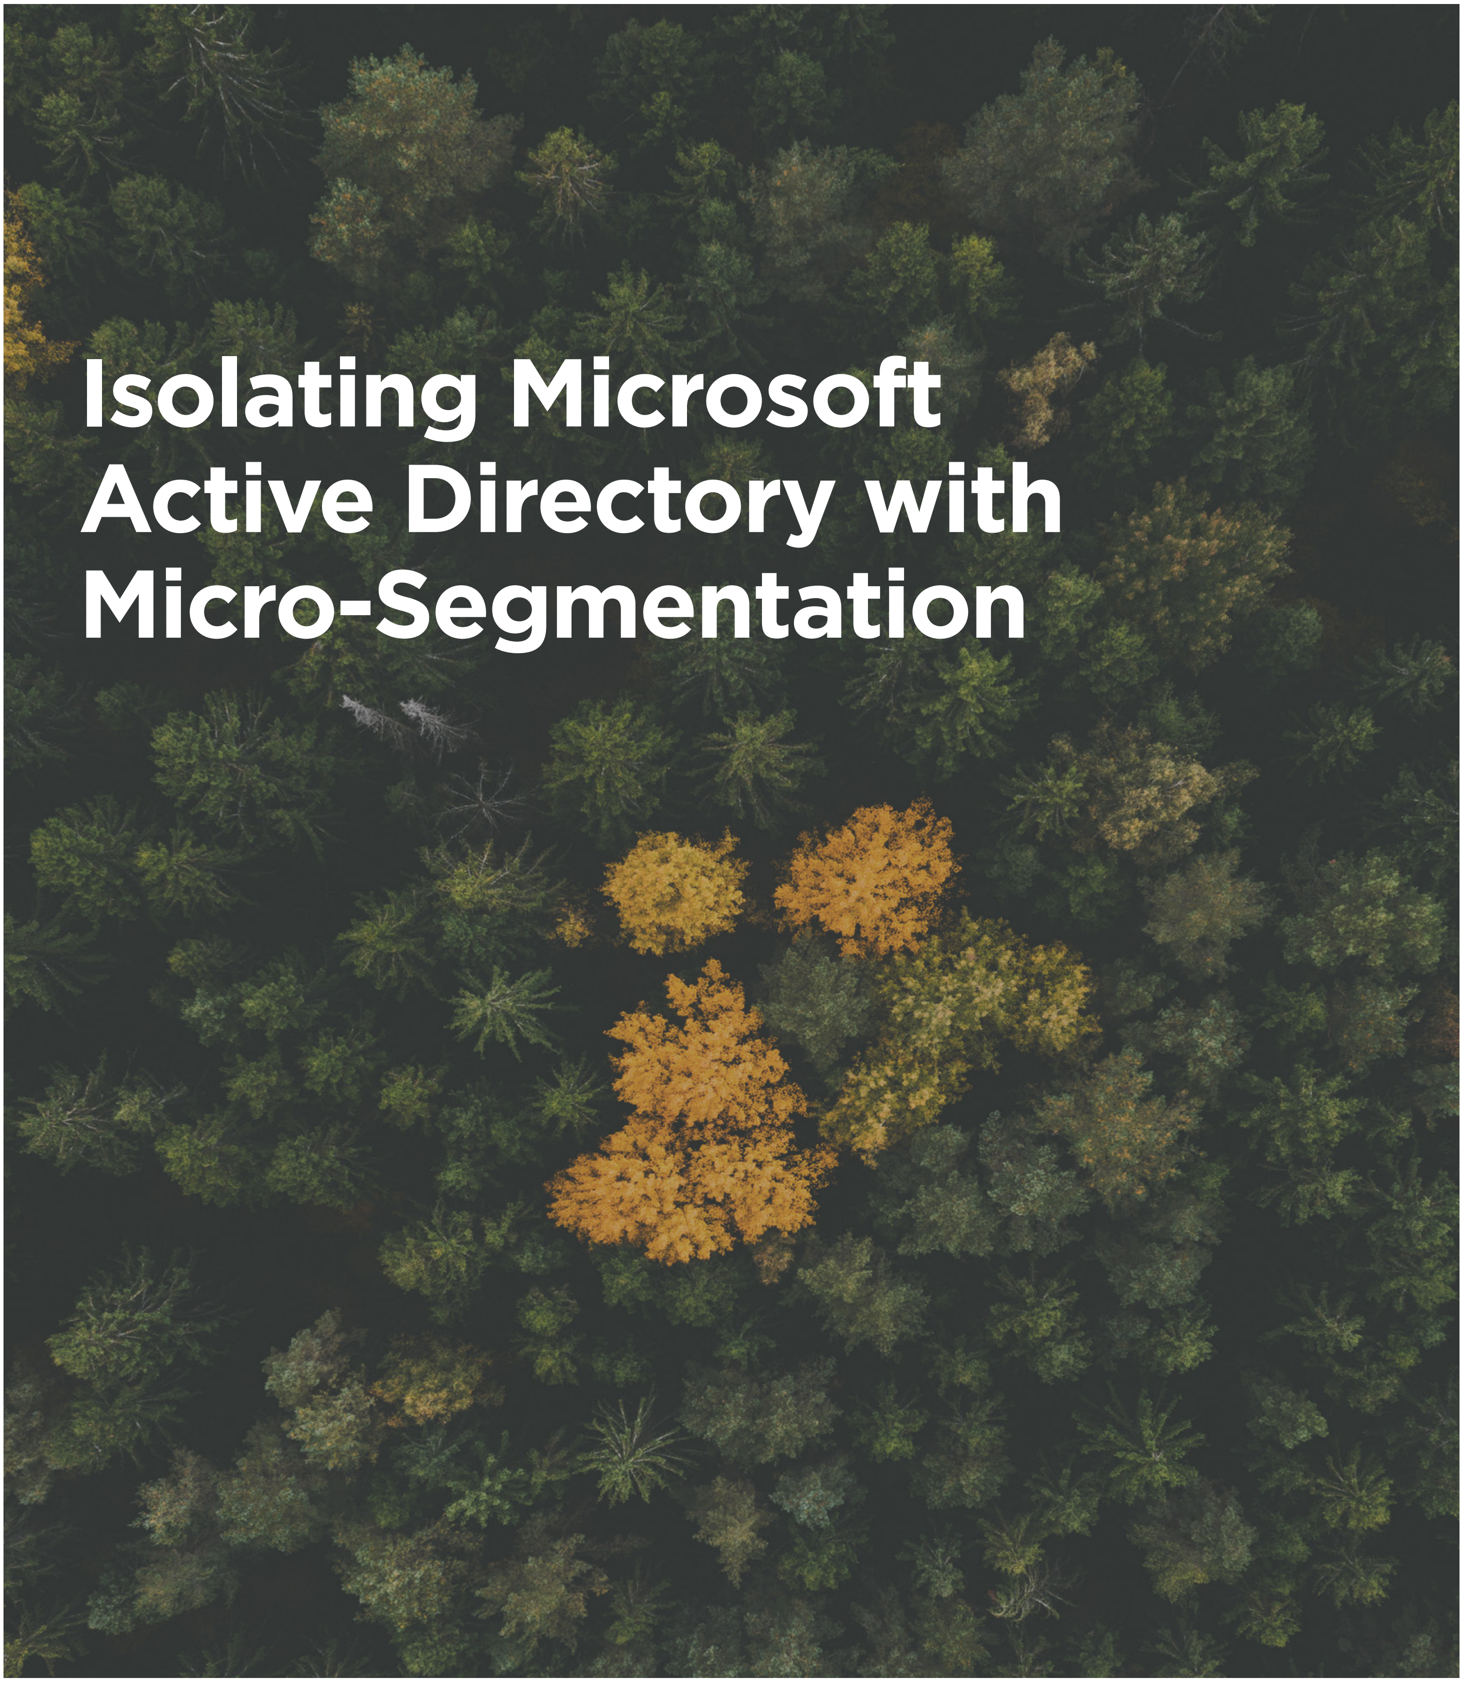

Overview

Active Directory Domain Controllers communicate with (and maintain information about) virtually every other server inside the environment. This role as central connective tissue makes Active Directory the perfect pivot point for an intruder looking to move laterally throughout data center and cloud environments.

This makes reducing Active Directory's attack surface an essential component of any security strategy for Microsoft environments. Unfortunately, many technologies aren't well-suited to address this challenge, forcing you to choose between an unacceptably large attack surface and an unwieldy security management burden that quickly breaks as your environment expands.

Contents

This white paper addresses:

Why is there an attack surface problem with Active Directory?

The attack surface of a server is all of the potential connections that another system in your environment (or on the Internet) could use to communicate with it. Within Active Directory, each potential connection consists of a process on one of your domain controllers that is actively listening on a given port, and another server with permission to communicate with it. The attack surface of Active Directory is all of the potential connections on domain controllers that an intruder could use to access it.

Attack Surface =

# of domain controllers ✖ **# of open ports on each domain controller** ✖ **# of systems (servers and clients) with permissions to communicate with Active Directory**

In other words, if one of your domain controllers has 10 ports open on it, and you have 50 servers in your environment with permission to communicate with that domain controller, then that domain controller would have an attack surface of 500 potential connections.

Each of these potential connections represents a doorway that an intruder can attempt to walk through. To be clear, the existence of a potential connection is not enough by itself – the intruder also needs a technique to exploit that connection. But these techniques are sadly easy to come by. WannaCry and NotPetya are two easy examples of recent pieces of dangerous malware that spread rapidly through networks and organizations by exploiting a simple vulnerability. In fact, both relied on a vulnerability in SMB – a process commonly found active on Active Directory domain servers.

Any potential connection could be vulnerable, which means if you have 500 potential connections, you have to secure all 500. If you only have 50 potential connections, you can concentrate your resources on this smaller attack surface.

As a result, isolating Active Directory (as well as other high-value applications) from unnecessary connections is a key element of any security strategy. This reduces the attack surface of not only your environment as a whole, but also any high-value applications within your environment.

Securing Active Directory with Network Segmentation

Unfortunately, securing your Active Directory servers requires managing far more than 500 ports. This is because, as any Windows Admin knows, Active Directory relies heavily on dynamic ports. The Active Directory process uses about 40 statically defined ports, but beyond that, it relies on a wide range of dynamic ports, both in high and low port ranges. In fact, an Active Directory installation could use as many as 20,386 dynamic ports, both in high (49,152-65,535) and low (1,024-5,000) port ranges.

DYNAMIC PORTS USED BY ACTIVE DIRECTORY

Low range

1,024 - 5,000

High range

49,152 - 65,535

Total used

20,386

Many organizations struggle to reduce the attack surface of their Active Directory deployment using network segmentation tools. The simple problem that they face is, while Active Directory can have many dynamic ports, it doesn't use them all the time. At any given moment, only a small fraction of these dynamic ports is actually in use.

Traditional firewall rules are keyed to static ports, which means that using traditional segmentation requires you to open all 20,386 ports, even if only a few of them are in use at any given time. This means that there are more than 20,000 ways an intruder could target each of your domain controllers, and chances are the vast majority of this exposure doesn't actually need to be open for your Active Directory deployment to function.

Isolating Active Directory with traditional firewalls requires that you either add domain controllers to each network segment and then configure them to talk to each other; or allow your clients and/or member servers to authenticate across network segments using the entire dynamic port range. Both of these solutions are problematic:

- Allowing port-based authentication across network segments radically expands the attack surface of Active Directory because you must open up those thousands of potential connections that are rarely (if ever) used for legitimate purposes.
- On the other hand, managing domain controllers for each segment makes your environment much more complex – increasing cost and increasing risk. Further, even if you reduce the cross-segment traffic to your domain controllers, you will still need to open large port ranges to enable communication.

Attack surface with network segmentation

Isolating Active Directory with network segmentation requires that you open all dynamic ports (20,386) on each domain controller (2) to communicate with all clients and servers in the environment (2,000). If this sounds like a large attack surface, it is. At best, you would have 81,544,000 allowed potential connections. Simply accessing an allowed connection isn't enough to exploit a device, but it does make it possible for an intruder to move laterally to access that device. That's 81,544,000 connections that your security team needs to worry about and secure.

Attack surface of Active Directory with network firewalls

2 domain controllers × **20,386** dynamic ports × **2,000** clients = **81,544,000** exposed connections

Host-Based Solutions

Network-based enforcement is simply not well-suited to the dynamic nature of Microsoft applications. This is why Microsoft recommends using host-based firewalls to isolate Active Directory and other key applications.

Host-based firewalls like the Windows Filtering Platform let organizations write dynamic, process-based security policies, which are well-suited to shrinking the exposure of applications like Active Directory. In fact, Microsoft itself recommends that organizations use host-based firewalls to isolate Microsoft applications (see Appendix A for more details).

Unfortunately, managing host-based firewalls at scale across a large organization creates significant challenges.

Some organizations attempt to manage this challenge by using firewall Group Policy Objects (GPOs). GPOs are essentially a collection of settings that define how a particular system will function for a defined group of users. This can help with some challenges, but setting unique policy for particular controllers or groups of controllers quickly becomes complex and cumbersome. Further, you still need to manually maintain links between IP addresses and system functionality. This can become challenging even with moderately-sized deployments of Active Directory, especially as you seek to apply more granular segmentation to reduce your attack surface.

Some more advanced enterprises have also tried modifying the registry to reduce the range of dynamic high ports used by their applications, and thus reduce their attack surface. While this can limit the range of dynamic ports in use (the RPC port mapper reads the registry for the range it can use), it also slows down authentication and forces you to maintain registry edits across all domain controllers. Out-of-sync systems can become unstable or fail.

Security teams really need a different approach that combines the granular control of host-based segmentation with the centralized coordination of more traditional firewall approaches.

Isolating Active Directory with Adaptive Micro-Segmentation

Micro-segmentation provides a powerful balance between both of these approaches, because it combines host-based sensors and points of enforcement with a centralized coordination platform. Host-based enforcement means that you can use granular, tailored enforcement to achieve just the right level of isolation around your most valuable systems. Centralized coordination means that you can avoid the problems with scale that host-based enforcement often carries.

Adaptive micro-segmentation takes this even further by deploying segmentation policies that can adapt as your environment changes. A perfect example of this are process-based policies. Micro-segmentation solutions that leverage Windows Filtering Platform can enforce process-based security policy as Microsoft recommends, allowing you to open only the minimum number of ports necessary at any given time.

Similarly, adaptive micro-segmentation ensures that as your environment shifts, the levels of isolation that you impose remain constant. For example, if you add a new domain controller to your Active Directory deployment, adaptive micro-segmentation will adjust to automatically absorb that new controller within its security graph, and ensure that it receives the correct level of isolation as soon as it spins up.

Attack surface with micro-segmentation

Because adaptive micro-segmentation can apply process-based security policies, you get a drastic reduction in your attack surface by reducing huge open port ranges down to just the dynamic processes that Active Directory requires. An Active Directory Domain Server could use any of around 20,000 ports – but it only uses about 20 processes. In our example of having 2 domain controllers and 2,000 clients and servers in that domain, you could reduce the 81,344,000 allowed port and protocol connections to Active Directory down to 80,000 (mainly process-based) connections using our pre-defined Active Directory Segmentation Templates. That's a reduction of more than 99.9 percent of the total attack surface of this particular Active Directory deployment.

Attack surface of Active Directory using Illumio Segmentation Templates

2 domain controllers × **20** processes × **2,000** clients = **80,000** exposed connections

99.9% reduction of total attack surface

What to Look for in a Micro-Segmentation Solution

Not all micro-segmentation solutions are created equal, and picking the wrong one can leave you struggling with dynamic environments, unable to apply Microsoft recommended process-based rules, and lagging when trying to operate at scale. If you're using a micro-segmentation solution to protect Active Directory (or any Microsoft application), it should be able to:

1. Enforce process-based policies

As discussed above, basing firewall rules on process, instead of port and protocol, actively ties network communications to a process on a workload. Instead of just allowing or denying a specific port and protocol, process-based policies allow you to designate exactly what software can use your network. Not all micro-segmentation solutions can enforce process-based security policies, however – many are still tied to port-based rules. If you're looking to reduce the attack surface of your Active Directory application (or any other Microsoft application), it's obviously essential that the micro-segmentation solution you choose can enforce process-based security policies.

2. Activate and manage the Windows Filtering Platform (WFP)

Some micro-segmentation platforms create their own enforcement – analyzing and blocking packets. But this can slow down your environment and leave you at the mercy of new and untrusted code as the solution develops. Even more important, all it does is reinvent the wheel. The Windows Filtering Platform, which is available on every modern Windows server, is already trusted, scalable, and fast – and it has built in process-based policy enforcement (this is the engine that GPO uses for those process-based rules that Microsoft recommends).

Using WFP gets you other advantages as well. Microsoft recognizes that processes and their associated files are targets, and has built significant capabilities into WFP to protect them. Windows doesn't just use SHA256 code signing certificates, but it also actively monitors them with the Windows Resource Protection Service. In other words, using WFP activates trusted security features that you've already paid for – making it the best option.

3. Offer quick-start resources like templates and policy-generating tools

One advantage of GPO is that you can use built-in Wizards to quickly deploy security policy on a host-by-host basis. Although this struggles at scale, the concept makes perfect sense. If you're looking at micro-segmentation solutions, look for one that gives you that same (or better) bootstrapping capability.

Something that GPO doesn't do is offer any "quick start" help if you're protecting a less standard application. But one of the key benefits that micro-segmentation offers is that by understanding what is happening on your domain controllers themselves, your security policy can understand what is happening in your environment as a whole, enabling it to dynamically identify and build security policy for even custom applications.

4. Adapt to changes in environments and threats by being application-aware

Environments today are dynamic – not static. If your segmentation solution is static and requires you to update it manually as your environment changes, it will never work at the speed and scale that you need.

For micro-segmentation to be adaptive, it must understand more than just your network – it must understand your applications. Segmentation that understands your applications can adapt as those applications change (expanding or contracting in response to demand, or shifting to new deployment platforms). Segmentation that only exists at the physical network layer or at the hypervisor can't do this, and so you will need to manually track and update your security policy as your environment changes. This leaves your team working constantly to keep up with your environment, and constantly at risk of making mistakes that could leave your systems exposed.

Application-aware micro-segmentation should also let you write policies in the language of your applications, rather than relying on traditional IP addresses. This makes your security policy much more readable and often much more concise. In most cases, security policy applied to applications can accomplish segmentation with less than 10 percent of the security rules required for IP-based policies.

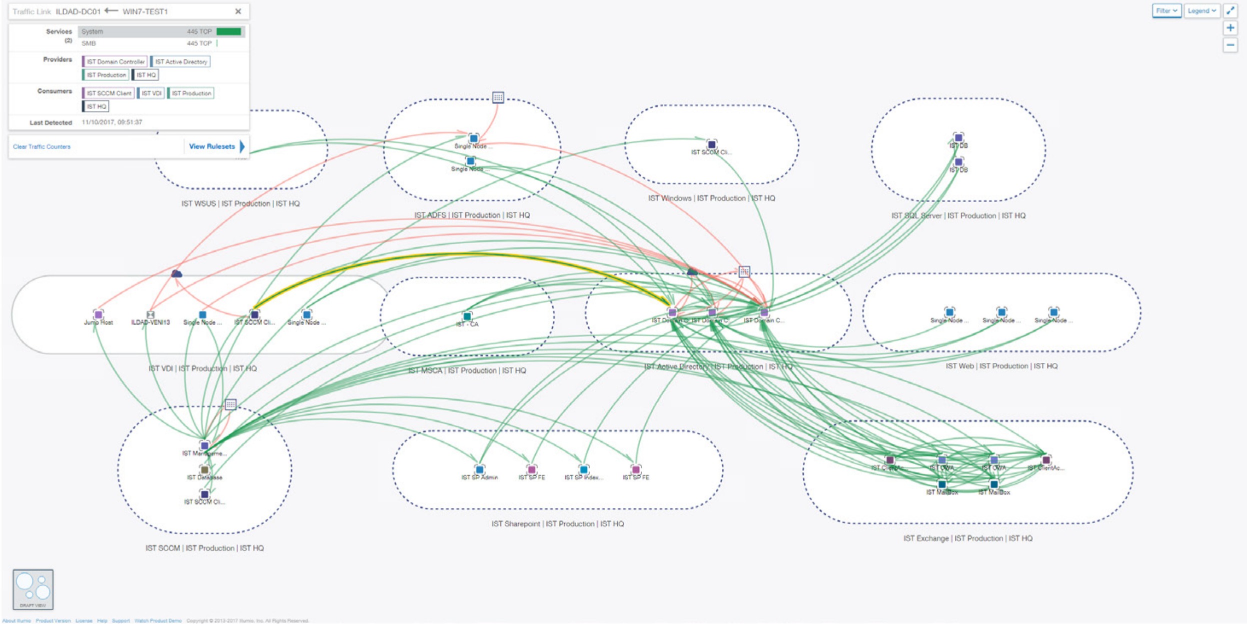

Traffic Link ILDAD-DC01 ← WIN7-TEST1

Services (2): System 445 TCP; SMB 445 TCP

Providers: IST Domain Controller, IST Active Directory, IST Production, IST HQ

Consumers: IST SCCM Client, IST VDI, IST Production, IST HQ

Last Detected: 11/10/2017, 09:51:37

Clear Traffic Counters
View Rulesets
Filter
Legend
IST WSUS | IST Production | IST HQ
Single Node
IST ADFS | IST Production | IST HQ
IST SCCM Cl...
IST Windows | IST Production | IST HQ
IST DB
IST SQL Server | IST Production | IST HQ
Jump Host
ILDAD-VEN13
IST VDI | IST Production | IST HQ
IST - CA
IST MSCA | IST Production | IST HQ
IST Active Directory | IST Production | IST HQ
IST Web | IST Production | IST HQ
IST Database
IST SCCM | IST Production | IST HQ
IST SP Admin
IST SP FE
IST SP Index...
IST Sharepoint | IST Production | IST HQ
IST Mailbox
IST Exchange | IST Production | IST HQ
DRAFT VIEW

An Active Directory environment – isolated using Illumio's adaptive segmentation.

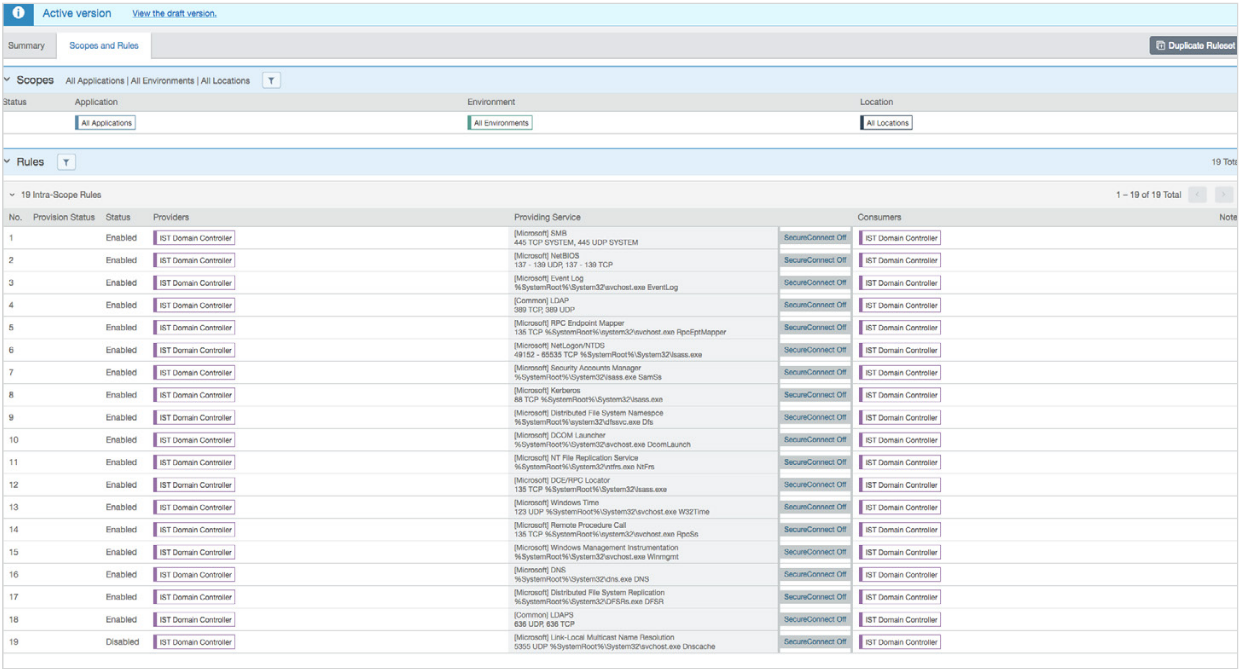

Scopes All Applications | All Environments | All Locations

Status	Application	Environment	Location
	All Applications	All Environments	All Locations

Rules — 19 Tot

19 Intra-Scope Rules — 1 – 19 of 19 Total

No.	Provision Status	Status	Providers	Providing Service		Consumers	Note
1		Enabled	IST Domain Controller	[Microsoft] SMB 445 TCP SYSTEM, 445 UDP SYSTEM	SecureConnect Off	IST Domain Controller	
2		Enabled	IST Domain Controller	[Microsoft] NetBIOS 137 - 139 UDP, 137 - 139 TCP	SecureConnect Off	IST Domain Controller	
3		Enabled	IST Domain Controller	[Microsoft] Event Log %SystemRoot%\System32\svchost.exe EventLog	SecureConnect Off	IST Domain Controller	
4		Enabled	IST Domain Controller	[Common] LDAP 389 TCP, 389 UDP	SecureConnect Off	IST Domain Controller	
5		Enabled	IST Domain Controller	[Microsoft] RPC Endpoint Mapper 135 TCP %SystemRoot%\system32\svchost.exe RpcEptMapper	SecureConnect Off	IST Domain Controller	
6		Enabled	IST Domain Controller	[Microsoft] NetLogon/NTDS 49152 - 65535 TCP %SystemRoot%\System32\lsass.exe	SecureConnect Off	IST Domain Controller	
7		Enabled	IST Domain Controller	[Microsoft] Security Accounts Manager %SystemRoot%\System32\lsass.exe SamSs	SecureConnect Off	IST Domain Controller	
8		Enabled	IST Domain Controller	[Microsoft] Kerberos 88 TCP %SystemRoot%\System32\lsass.exe	SecureConnect Off	IST Domain Controller	
9		Enabled	IST Domain Controller	[Microsoft] Distributed File System Namespce %SystemRoot%\system32\dfssvc.exe Dfs	SecureConnect Off	IST Domain Controller	
10		Enabled	IST Domain Controller	[Microsoft] DCOM Launcher %SystemRoot%\System32\svchost.exe DcomLaunch	SecureConnect Off	IST Domain Controller	
11		Enabled	IST Domain Controller	[Microsoft] NT File Replication Service %SystemRoot%\System32\ntfrs.exe NtFrs	SecureConnect Off	IST Domain Controller	
12		Enabled	IST Domain Controller	[Microsoft] DCE/RPC Locator 135 TCP %SystemRoot%\System32\lsass.exe	SecureConnect Off	IST Domain Controller	
13		Enabled	IST Domain Controller	[Microsoft] Windows Time 123 UDP %SystemRoot%\System32\svchost.exe W32Time	SecureConnect Off	IST Domain Controller	
14		Enabled	IST Domain Controller	[Microsoft] Remote Procedure Call 135 TCP %SystemRoot%\system32\svchost.exe RpcSs	SecureConnect Off	IST Domain Controller	
15		Enabled	IST Domain Controller	[Microsoft] Windows Management Instrumentation %SystemRoot%\System32\svchost.exe Winmgmt	SecureConnect Off	IST Domain Controller	
16		Enabled	IST Domain Controller	[Microsoft] DNS %SystemRoot%\System32\dns.exe DNS	SecureConnect Off	IST Domain Controller	
17		Enabled	IST Domain Controller	[Microsoft] Distributed File System Replication %SystemRoot%\System32\DFSRs.exe DFSR	SecureConnect Off	IST Domain Controller	
18		Enabled	IST Domain Controller	[Common] LDAPS 636 UDP, 636 TCP	SecureConnect Off	IST Domain Controller	
19		Disabled	IST Domain Controller	[Microsoft] Link-Local Multicast Name Resolution 5355 UDP %SystemRoot%\System32\svchost.exe Dnscache	SecureConnect Off	IST Domain Controller	

Illumio's Segmentation Templates let you implement many security policies with the click of a button.

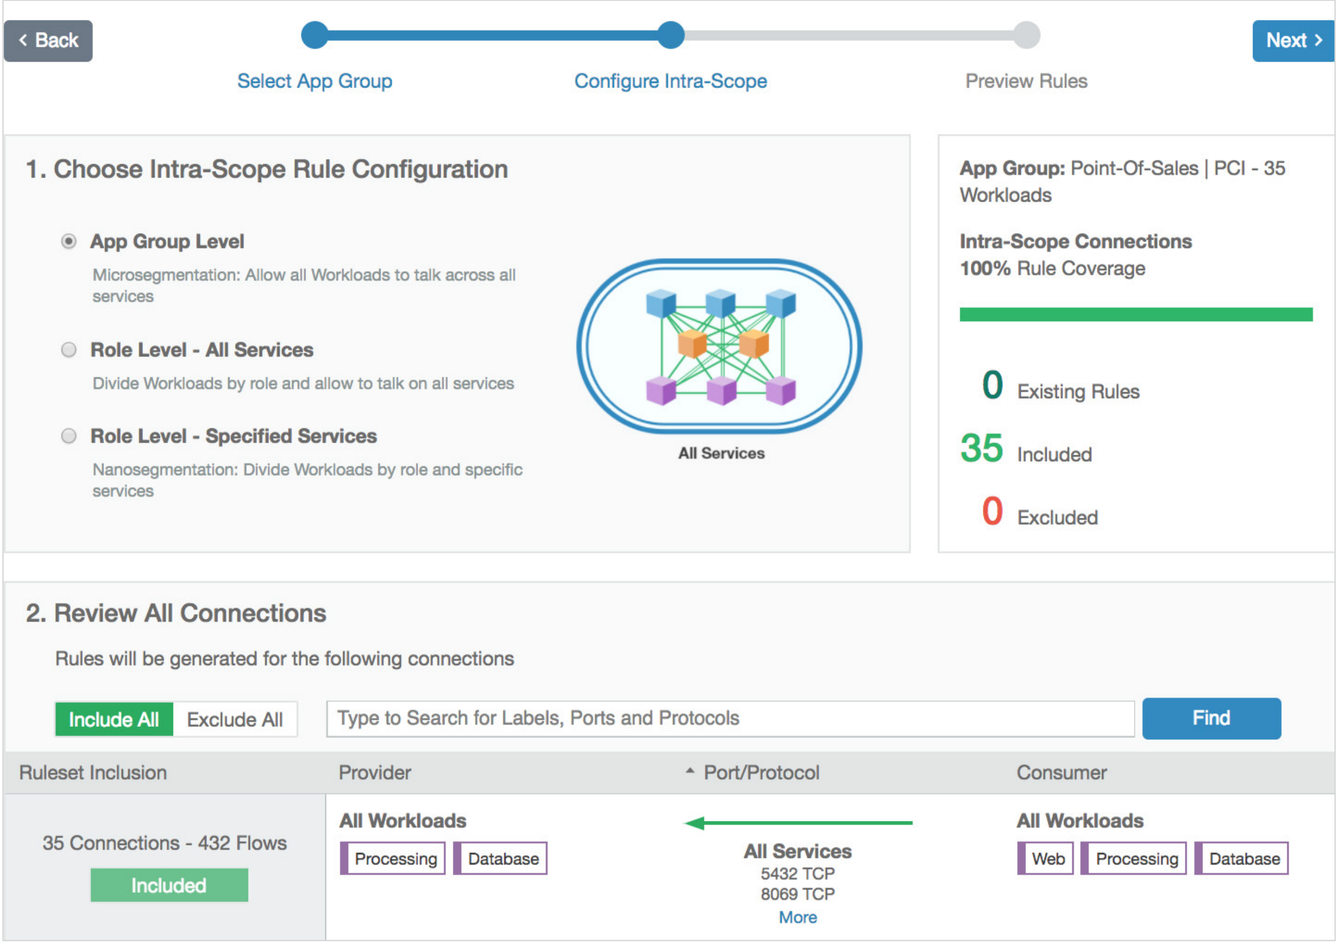

1. Choose Intra-Scope Rule Configuration

● App Group Level
Microsegmentation: Allow all Workloads to talk across all services

○ Role Level - All Services
Divide Workloads by role and allow to talk on all services

○ Role Level - Specified Services
Nanosegmentation: Divide Workloads by role and specific services

App Group: Point-Of-Sales | PCI - 35 Workloads

Intra-Scope Connections
100% Rule Coverage

0 Existing Rules
35 Included
0 Excluded

2. Review All Connections

Rules will be generated for the following connections

Type to Search for Labels, Ports and Protocols

Ruleset Inclusion	Provider	Port/Protocol	Consumer
35 Connections - 432 Flows Included	All Workloads Processing Database	All Services 5432 TCP 8069 TCP More	All Workloads Web Processing Database

Illumio's Policy Generator automatically generates security policies for custom applications.

Appendix

Best Practices for Securing Active Directory
https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/manage/component-updates/executive-summary

Securing Domain Controllers Against Attack
https://technet.microsoft.com/windows-server-docs/identity/ad-ds/plan/security-best-practices/securing-domain-controllers-against-attack

Securing Privileged Access
https://technet.microsoft.com/en-us/windows-server-docs/security/securing-privileged-access/securing-privileged-access

Illumio enables organizations to realize a future without high-profile breaches by preventing the lateral movement of attackers across any organization. Founded on the principle of least privilege in 2013, Illumio provides visibility and segmentation for endpoints, data centers or clouds. The world's leading organizations, including Morgan Stanley, BNP Paribas, Salesforce, and Oracle NetSuite, trust Illumio to reduce cyber risk. For more information, visit www.illumio.com/what-we-do.

Gartner peerinsights™

See what customers have to say about Illumio.

The GARTNER PEER INSIGHTS Logo is a trademark and service mark of Gartner, Inc. and/or its affiliates and is used herein with permission. All rights reserved. Gartner Peer Insights reviews constitute the subjective opinions of individual end users based on their own experiences and do not represent the views of Gartner or its affiliates.

Follow us on:

illumio